

UNITED STATES PATENT APPLICATION

OF

REINHARD DEML

and

ANDREAS STRAUB

FOR

**PERFORMANCE IMPROVEMENT FOR
ATM AAL2/5 TO IP PACKET PROCESSING**

**Attorney Docket No. 033387-002
BURNS, DOANE, SWECKER & MATHIS, L.L.P.
P.O. Box 1404
Alexandria, Virginia 22313-1404
(703) 836-6620**

09975087-10001
FOOT" 2805/560

PERFORMANCE IMPROVEMENT FOR ATM AAL2/5 TO IP PACKET PROCESSING

Background of the Invention

[0001] The present invention relates to network interfaces of universal mobile telecommunication systems (UMTS) radio network controllers. The network interface of a UMTS radio network controller routes asynchronous transfer mode (ATM) adaptation layer 2 (AAL2) and ATM adaptation layer 5 (AAL5) packetized data to an Ethernet-based Internet Protocol (IP) network.

[0002] Typically, these operations are implemented in software that uses parameters stored in lookup tables and creates the protocol headers from scratch in order to encapsulate the AAL2/AAL5 packets in an Ethernet Frame. The Ethernet header, IP header, User Datagram Protocol (UDP) packet header, and Ethernet frame Cyclical Redundancy Check (CRC) information are produced. This results in a significant amount of calculations for each Ethernet frame produced. These calculations can significantly slow the operation of the system.

[0003] It is desired to have a system with improved performance for implementing the encapsulation of AAL2/AAL5 packetized data in a network interface of a UMTS radio network controller.

Summary of the Present Invention

[0004] In the present invention, partial header data is buffered for each session. Each AAL2/AAL5 packet that belongs to a particular session is associated with the same predefined partial header data stored in the buffer. The partial header data does not include all of the fields of the header. The partial header data along with the AAL2/AAL5 packet is provided to hardware which calculates the length information and error check information (checksums and CRC) required to produce a complete Ethernet frame.

0975087-100501

[0005] By setting up the header buffer for each session and by using hardware to calculate portions of the header and the Ethernet frame CRC, the speed of the system of the present invention is significantly improved. The performance gain of one embodiment of the system of the present invention compared to an off-the-shelf network protocol stack is about eight times. Even if the user optimizes the prior-art protocol stack, the performance gain of one embodiment of the system of the present invention is in the range of 4-5 times.

[0006] The ATM header of the incoming AAL2/AAL5 packets is used to indicate the session and thus indicate which stored partial header information is provided to the hardware.

[0007] In a preferred embodiment, the partial header information is calculated once per session and stored in a data buffer. The linked list of data buffers includes a pointer to the partial header data for that session and a pointer to the buffer storing the AAL2/AAL5 packet.

[0008] In one embodiment of the present invention, a field in incoming packets is used to access stored partial header information. In hardware, the stored partial header information and the incoming packets are used to calculate additional information. The additional information including at least one length field and at least one error check field, the output of the hardware being the incoming packet encapsulated by one or more protocols.

[0009] Another embodiment of the present invention is a method for each session, constructing and storing partial header information. The partial header information including source and destination fields. Using a field in the incoming packets to access stored partial header information, the same partial header information being used for each incoming packet of the session. In hardware, using the stored partial header information and the incoming packets to calculate additional information. The additional information including at least one length

09975067-100901

field and at least one error check field. The output of the hardware being one of the incoming packets encapsulated by one or more protocols.

[0010] In one embodiment, the protocols include the UDP protocol, the IP protocol, and an Ethernet protocol. In one embodiment the additional information includes the UDP message length, a UDP checksum value, an IP header checksum, an IP total length value, an Ethernet frame payload length, and an Ethernet frame CRC value.

[0011] One embodiment of the present invention is a UMTS system including a radio network controller including a network interface. The network interface includes software adapted to store partial header information for a session. The session indicated by a field of incoming packets. The stored partial header information including source and destination information for at least one protocol. The network interface also including hardware receiving the stored partial header information in the incoming packets. The hardware is adapted to calculate additional information for the outgoing data. The additional information including at least one length field and at least one error check field.

[0012] In one embodiment, a data transmitter machine according to the invention is built in two stages, which are pipelined. The first part of the machine produces data packets represented by a sequence of a header and the data block. The header information transferred between the two stages contains some data fields which are not evaluated to their final value. The data block contains a predefined number of parameters(data fields), which are inserted into the data block to the positions transferred immediately after the header. The second stage uses the parameter transferred in the data block to compute the final value of the header data. The first stage may be a conventional interface component with software providing the half-evaluated header blocks and linking the data blocks with the inclusion of the individual parameter information to the head of the data block. The second stage is typically hardware, producing the correct header computed from the half-evaluated

099503-100904

header and the parameter informations, and then transferring the data block (without the parameters).

Brief Description of the Drawings

[0013] Fig. 1 illustrates a UMTS wireless network which can use the system of the present invention.

[0014] Fig. 2 is a diagram that illustrates the AAL2/AAL5 packetized data encapsulation of one embodiment of the present invention.

[0015] Fig. 3 is a diagram that illustrates one embodiment of the present invention using a software unit and a hardware unit.

[0016] Fig. 4 is a diagram that illustrates the hardware and software unit contributions to the Ethernet frame encapsulation of one embodiment of the system of the present invention.

[0017] Fig. 5 is a diagram that illustrates the header information used in one embodiment of the present invention.

[0018] Fig. 6 is a diagram that illustrates a linked data buffer used in one embodiment of the present invention.

[0019] Fig. 7 is a diagram that illustrates an AAL2/AAL5 packet with an appended length value which is used in one embodiment of the system of the present invention.

[0020] Fig. 8 is a diagram that illustrates a hardware unit of one embodiment of the system of the present invention.

[0021] Fig. 9 is a diagram that illustrates details of one embodiment of a hardware unit of the system of the present invention.

Detailed Description of the Invention

[0022] Fig. 1 is a diagram of a UMTS wireless network. The network includes a core network 22, a radio access network 24 and wireless devices 26 and 28. The

FOOTNOTES

radio access network 24 includes a Node B 30 and a radio network controller 32. The radio network controller 32 includes a network interface 34. In one embodiment the network interface receives ATM cells and reassembles a AAL2/AAL5 packet and encapsulates these packets in additional protocols. In a preferred embodiment, the AAL2/AAL5 packets are encapsulated in the UDP protocol, the IP protocol, and an Ethernet protocol.

[0023] Fig. 2 illustrates a network interface of one embodiment of the system in the present invention. In this embodiment, ATM cells are received, reassembled and then, based upon the ATM cell header of AAL2/AAL5 packets, are routed using the Ethernet protocol. In one embodiment, AAL2 packets are routed based upon the ATM Virtual Channel (VC) and the sub-channel identifier (CID); the AAL5 packets are routed based upon the ATM virtual channel (VC). Each ATM AAL5 connection and ATM AAL2 sub-channel is mapped to a unique UDP/IP session identified by an IP destination address and a UDP destination port number. The VC (for AAL5 packets) or VC and CID (for AAL2 packets) are used to access a routing table which allows the production of the header and other overhead associated with the protocols.

[0024] An ATM cell header consists of an 8 bit virtual path identifier and a 12 bit virtual channel identifier. Together they identify the ATM virtual channel (VC). For AAL5 packets, only the ATM cell header, VC, is used to identify the session. The VC information can be added as additional information inside the AAL5 packet. For AAL2 packets, the VC from the ATM cell header and the subchannel ID (CID) which is part of the AAL2 packets are used to identify the session.

[0025] In prior-art systems, this overhead calculation is done purely in software. The network driver takes the parameters and creates the protocol header and footer required. This can be quite time-consuming since it requires a lot of calculations and memory transactions.

[0026] Fig. 3 illustrates a system of one embodiment of the present invention. In this embodiment, the ATM cells are received at an ATM reassembly unit 40. The ATM reassemble unit takes the ATM cells and reconstructs AAL2/5 packets. In a preferred embodiment, the AAL2/AAL5 packet has appended to it a packet length value. The packets are sent to software unit 42. In a preferred embodiment, the software unit, for each IP session, prepares a partial header and buffers it in the buffer 44 of the memory 46. As will be described below, the buffer is preferably set up as a linked data buffer 44. In one embodiment, each of the AAL2/AAL5 packets are buffered in the linked data buffer 44. The software unit preferably maintains the linked data buffer 44.

[0027] The partial header information set up for each IP session preferably includes the source and destination information for each of the protocols used. The partial header information and the AAL2/AAL5 packet is provided to the hardware unit 50. Hardware unit 50 calculates the additional information required and updates the header and, in a preferred embodiment, calculates the Ethernet frame CRC footer. Thus, the output of the hardware unit 50 is an Ethernet frame which can be sent to an IP switch 52, which switches based upon the IP protocol.

[0028] Fig. 4 is a diagram that illustrates the contributions from the hardware and the software to the constructed Ethernet frame in one embodiment of the system of the present invention. The AAL2/AAL5 packet acts as a payload for UDP packet which includes a UDP header and the payload. The UDP packet is a payload for the IP packet which further includes an IP packet header. The IP packet acts as the payload for the Ethernet frame which further adds the Ethernet frame header and the Ethernet frame CRC footer. The header is comprised of the Ethernet frame header, the IP packet header and the UDP header. These protocols are used in a preferred embodiment, but other protocols can be used. Partial header information is provided from the buffer. This partial header information includes source and destination information such as the source and destination

address and source and destination port. Additional header information is produced in the hardware. In a preferred embodiment, the Ethernet frame CRC is also calculated in hardware.

[0029] Fig. 5 illustrates the header information of one embodiment of the system in the present invention. The header information includes an Ethernet frame header portion, an IP header portion and a UDP header portion. The Ethernet frame header portion includes a destination Medium Access Control (MAC) address, a source MAC address and the Ethernet type field. The source and destination addresses are determined in software for once for each session. In the example of Fig. 3, the source MAC address is the MAC address of unit 51. The destination MAC address is the MAC address of the IP switch 52. The Ethernet type for payloads of 1536 or fewer bytes indicates the length of the Ethernet frame payload. As will be described below, this length is preferably calculated in hardware.

[0030] The IP header portion includes a version field, a type of service field, identification field, a fragmentation field, a time to live field, protocol ID field, the source IP address and destination IP address. Many of the fields are fixed for each IP session. The source and destination IP address and UDP source port and destination port depend upon a field in the AAL2/AAL5 packet which provides an index to a routing table. Note that the IP header checksum and total length are preferably calculated in hardware. The UDP header portion includes a UDP source port, a UDP destination port which are determined in software and UDP message length and UDP checksum which are calculated in hardware.

[0031] The partial header information is set up once per IP session when the IP session is established and reused for each incoming packet of the IP session. Instead of creating each IP packet from scratch as is done in the prior art network software drivers, the data buffers only have to be linked together with transmit

buffer descriptors. The software just links but does not modify the partial header data buffer after the partial header data buffer is constructed.

[0032] Looking again at Fig. 3, in one embodiment, the routing table 45 uses a field of the AAL2/AAL5 incoming packets, such as the VC or a CID field, as an index to determine the partial header pointer. If a new session is being set up, the partial header needs to be constructed as well as a pointer to the partial header. The pointer information is added to the routing table 45.

[0033] Fig. 6 is a diagram that illustrates the linked data buffer of one embodiment of the present invention. As each packet is received, a field of the incoming packet is checked in a routing table to determine a partial header pointer. If there is a partial header pointer, this pointer is added as the first buffer descriptor for the packet. If there is not a partial header pointer, the VC (for AAL5 packets) or VC and CID (for AAL2 packets) indicates a new session, and the software unit will construct the partial header information. The partial header information indicates the source and destination for the different protocols. The partial header is stored in the buffer. Also a pointer to this buffer location will also be produced and added to the routing table 45 associated with the VC or VC and CID. Note that each AAL2/AAL5 packet received for a session has its first buffer descriptor pointing to the same partial header data. The other buffer descriptors of a packet point to the buffer locations where the AAL2/AAL5 packet is stored.

[0034] For each AAL2/AAL5 packet, the linked list will include a pointer to the covered the partial header information as well as a pointer to the packet. The hardware uses a hardware read pointer to obtain the pointers to the partial header information and the packet. The linked list data buffer 42 transfers the incoming packets in a first-in first-out fashion.

[0035] Fig. 7 illustrates the AAL2/AAL5 packet with the appended length information. The appended length information is determined in the AAL ATM

reassembly process. By appending this length information to the packet, the operations of the hardware unit 50 can be simplified. Note that in other embodiments the software or hardware unit itself calculates the length, but it is simplest to append the length in the reassembly unit.

[0036] Fig. 8 is a diagram of one embodiment of a hardware unit of one embodiment of the system of the present invention. The hardware unit receives the incoming packet, in one embodiment the AAL2/AAL5 packet with appended packet length. A packet length determination unit 62 separates the packet length from the incoming packet. The AAL2/AAL5 packet length information is used to calculate the length fields for the different protocols. In this embodiment the UDP, IP and Ethernet protocols are used. The partial header includes a partial UDP header provided to the UDP header hardware unit 64, a partial IP header provided to the IP header hardware unit 66 and a partial Ethernet header provided to the Ethernet header hardware insert unit 68. The UDP header hardware inserts include a UDP packet length insert and a UDP packet checksum insert. The UDP packet length can be calculated by adding the UDP header length, which is a constant, to the AAL2/AAL5 packet length. The UDP packet checksum is a checksum of the UDP packet including the header and the payload, in this case the AAL2/AAL5 packet. The IP header hardware unit 66 determines the IP packet length which is, in this case, the AAL2/AAL5 packet length added to a constant for the combined IP header and UDP header length. The IP header checksum is a checksum on the IP header excluding the payload, meaning that the unit 66 does not depend upon the calculations of the unit 64 and can operate in parallel. The Ethernet header hardware unit 68 determines the Ethernet payload length. As noted above, for payloads of 1536 bytes or less, the Ethernet type field contains the Ethernet payload length which, in this case, is the same as the IP packet length. The outputs of the units 64, 66 and 68 include the complete header. The header information along with the AAL2/AAL5 packet is provided to the Ethernet

FOUO 48054660

frame cyclical redundancy check (CRC) unit 70. Note that the operations of the Ethernet frame cyclical redundancy check unit 70 is dependent upon the operation of the header insert units 64, 66, and 68 and must wait for these units to operate it before calculating the CRC.

[0037] A cyclical redundancy check for a k-bit input determines an n-bit sequence known as a frame check sequence so that the resulting frame consisting of $k + n$ bits is exactly divisible by a predetermined number. This number is indicated by the divisor polynomial, $P(X)$. In the one embodiment, the CRC logic is implemented as a dividing circuit consisting of exclusive-or gates input into a shift register. The exclusive-or gate arrangement is determined by the divisor polynomial.

[0038] Fig. 9 illustrates details of one embodiment of the hardware unit 72 of one embodiment of the present invention. In this embodiment, note that the UDP packet length is calculated by an adder 74 adding the incoming packet length to the UDP header length constant. The UDP checksum calculation unit 76 uses the updated UDP header information along with the incoming packet to calculate the UDP checksum value which is then inserted to produce the complete UDP header. The adder 78 adds the incoming packet length to the combined IP and UDP header length constant to produce the IP packet length. Note that the IP packet length is also the Ethernet payload length which is inserted into the partial Ethernet header. The IP checksum unit 80 determines the checksum for the IP header using the partial IP header after the IP packet length is inserted. As discussed above, the output of the adder 78 can be used in both the IP header hardware insert unit and the Ethernet header hardware insert unit. Alternately, multiple adder units can be used.

[0039] Note that the embodiments of Figs. 8 and 9 are merely one example of a hardware unit and a number of arrangements can be used to implement the hardware unit. In a preferred embodiment the hardware unit is implemented in a

T06007-10090

field programmable gate array. Alternately the hardware unit can be implemented in any other fashion such as in an application specific integrated circuit (ASIC).

[0040] It will be appreciated by those of ordinary skill in the art that the invention can be implemented in other specific forms without departing from the spirit or character thereof. The presently disclosed embodiments are therefore considered in all respects to be illustrative and not restrictive. The scope of the invention is illustrated by the appended claims rather than the foregoing description, and all changes that come within the meaning and range of equivalents thereof are intended to be embraced herein.

10595037-10001